REMARKS

The present response is filed in response to the Office Action of April 25, 2003.

Claims 1-16 are presently pending in the application and claims 1-16 have been rejected under 35 U.S.C. § 102 and under 35 U.S.C. § 103 over U.S. Patent No. 6,205,449 to Rastogi et al. for a SYSTEM AND METHOD FOR PROVIDING HOT SPARE REDUNCANCE AND RECOVERY FOR A VERY LARGE DATABASE MANAGEMENT SYSTEM, hereafter referred to as "Rastogi et al. `449".

More specifically, claims 1, 3, 5, 9, 11, 13 and 15 are rejected over Rastogi et al. `449 under 35 U.S.C. § 102(e) while claims 2, 4, 6, 8, 10, 12, 14 and 16 are rejected over Rastogi et al. `449 under 35 U.S.C. § 103(a) The Applicant acknowledges and respectfully traverses the raised obviousness rejection in view of the following remarks.

The Examiner has also objected to the Abstract of the specification as exceeding the allowed word count for an Abstract.

First, the Applicant has responded to the Examiner's objection to the Abstract by providing a new Abstract based on the previous Abstract but within the required word count for an Abstract. The Applicant, therefore, respectfully requests that the Examiner enter the Abstract submitted herein and that the Examiner reconsider and withdraw the objection to the Abstract.

Next considering the rejections of the claims under 35 U.S.C. § 102 and 35 U.S.C. § 103, the present invention and the claims of the present application are directed to system resource for performing system resource operations requested by a client of the system resource, such as a file server performing file transactions requested by a client, wherein the system resource includes a state logging mechanism for extracting, storing and restoring information representing the state of operation of the system resource and of the system resource operations, such as file transactions. The present invention and the claims of the present Application are directed to the method of operation of such a system resource, to the state logging mechanism and to the method of operation of the state logging mechanism.

In a first aspect of the present invention, the claims recite and define a system resource, such as a file server, with a state logging mechanism as including a system resource sub-system, such as a storage sub-system, and a control/processing sub-system that includes a resource control processor performing system resource operations in response to client requests and controlling operations of the system resource sub-system, such as a file system processor performing file transaction operations in response to client requests and controlling file storage operations of the storage sub-system, and a state machine logging mechanism. The state machine logging mechanism, in turn, includes a state machine log generator for extracting state machine information defining a state machine representing a current state of execution of a system resource operation and a state machine log mechanism for storing the state machine information. The state machine log generator is responsive to the restoration of operation of the system resource after a failure of system resource operations for reading the state machine information from the state machine log mechanism and restoring the state of execution of a system resource operation.

As described in the specification and as recited in the claims, the state machine log mechanism includes a state machine log mirroring mechanism located separately from the control/processing sub-system and communicating with the state machine log generator for receiving and storing mirror copies of the state machine information, wherein the state machine log mirroring mechanism is responsive to the restoration of operation of the file server after a failure of file server operations for reading the state machine information from the state machine log mirroring mechanism and restoring the state of execution of a file transaction.

Also according to the present invention, and as described and defined in the specification, "state machine information" is defined and described as and by the control and data values residing in the machine during a given state, or point in time.

A "state machine", in turn, is described and defined in the specification as a machine or system, such as a system resource or a file server, that executes operations as a sequence of discrete operating "states" wherein, as discussed, a "state" is defined and described by the

control and data values residing in the machine during that state, or point in time. The present and next operating state of a state machine are described and defined by the current state of the machine and the state functions of the machine itself, that is, the logic and circuit functions implemented in the machine that determine the responses or changes in state of the machine as a result of a current operating state. A "state machine", such as a system, sub-system or logical or functional element of a system or sub-system of any form, may therefore may be defined and described as a sequence of state machines wherein each state machine in the sequence of state machines is defined by the current state, that is, control and data values residing in the machine, and the state functions of the machine, that is, the functions or operations that will be executed by the state machine to result in the next state machine. It will be apparent that, for a given system, the state functions of the state machines describing and defining the system, are fixed and implicitly known and need not be specified for each state machine individually. As such, the current state of operation of a system may be defined by the state of the current state machine, that is, the control and data values residing in the state machine, and a sequence of operations executed by the given system may be defined and described by the corresponding sequence of states of the state machines.

At this point, it must be noted that the state machine logging mechanism of the present invention extracts and stores "state machine information" defining one or more "state machines" during the execution a resource operation or transaction rather than only at the conclusion of the operation or transaction, as is done in transaction logging type systems.

The state machine logging mechanism of the present invention thereby provides greater "granularity" in recording the execution of an operation or transaction in allowing the capture and storing of a sequence of state machines during the performance or execution of a resource operation as transaction, each state machine being comprised of the state machine information at each of a sequence of points in the performance or execution of the operation or transaction information. The state machine logging mechanism of the present invention thereby allows the

restoration of a resource server or transaction at a point during the performance or execution of the operation or transaction.

The state machine logging mechanism of the present invention thereby also provides greater assurance that a given resource server operation or transaction is captured for subsequent restoration, if needed, because the necessary information is captured at any of a sequence of points in the execution of the operation or transaction, rather than only when the operation or transaction is completed.

Lastly, it must be noted that the invention as described in the specification and as recited in certain of the claims is directed to an implementation of the present invention in a system resource having a single control/processing sub-system having an associated state machine log generator and a state machine log, such as in claims 1, 3, 5, 7, 9, 11, 13 and 15, and to an implementation wherein the state machine logging mechanism further includes a state machine log mirroring mechanism that is separate and independent form the control/processing sub-system, such as in claims 2, 4, 6, 8, 10, 12 14 and 16. That is, in the recitations of claims 2, 4, 6,8, 10, 12, 14 and 16, the state machine logging mechanism has a state machine log situated in or in association with the control/processing sub-system and a separate state machine log mirroring mechanism that is effectively external and independent of the control/processing sub-system.

In this regard, it must be noted that, as described in the specification, the state machine log mirroring mechanism is not a separate system analogous to the control/processing sub-system, but is a specific purpose mechanism designed and constructed only to perform the state machine log mirroring function of capturing and storing a copy of the state machine information and state machines that are stored in the state machine log associate with the state machine log generator that is associated directly with the control/processor sub-system. As described, the state machine log mirroring mechanism is separate and independent from the corresponding control/processing with respect to such matters as power supplies and similar elements and functions that perform support functions rather than being directly involved

in the execution of resource server operations or transactions and the logging of operation and transactions so that a failure of the corresponding control/processing sub-system itself will not cause a failure in the state machine log mirroring mechanism.

It must also be noted that the claims are directed to implementations of the present invention having a single control/processing sub-system with an associated state machine logging mechanism and possibly having an associated state machine log mirroring mechanism, depending on the specific implementation and claim, and to implementations having dual control/processing sub-systems with associated state machine logging mechanisms and state machine log mirroring mechanisms. In these implementations, each control/processor sub-system has an associated state machine logging mechanism for its own transactions or operations and a state machine log mirroring mechanism associated with the other control/processing sub-system. As such, each control/processing sub-system maintains its own state machine logging mechanism and maintains a state machine log mirroring mechanism for the other control/processing sub-system.

It must also be noted that in the system resources or file servers having dual control/processing sub-systems, the two control/processing sub-systems operate concurrently, independently and in parallel with each other and with each control/processing sub-system performing its own system resource operations or transactions, and with each control/processing sub-system providing only a residence and support for the state machine log mirroring mechanism of the other control/processing sub-system.

As such, and as described in the specification and recited in the claims, neither control/processing sub-system provides "back-up" for the other control/processing system in the sense of performing or executing transactions or operations for requests directed to the other control/processing sub-system. It must also be noted that because of the concurrent, independent and parallel operation of each control/processing sub-system, with each control/processing sub-system executing only the requests for transactions or operations that are directed to it, the full processing power of the two control/processing sub-systems is

typically available at all times to process requests from clients. If one control/processing sub-system fails, the other control/processing sub-system remains operative to perform the request directed to it, and to maintain the state machine information necessary to subsequently restore the failed control-processing sub-system. In addition, the operative one of the control/processing sub-systems may perform requests that would be directed to the other control/processing sub-system if those requests are re-directed to the operative control/processing sub-system.

In this regard, it must be noted that in conventional "back-up" systems there typically two identical systems arranged in parallel but, because one unit is a 'back-up" for the other, the "back-up" unit is typically idle until it is required to replace the other unit in performing the operations, so that only one half of the total power and resources of the system are available at any time. The system of the present invention, however, makes significantly more of the potential control/processing sub-system available to clients of the system.

Next considering the teachings of Rastogi et al. '449, Rastogi et al. '449 describes a system having "hot spare" support wherein the system includes two identical computer systems communicating through a network and wherein, at any given time, one of the computer systems is designated as the "primary" computer system and the other is designated as the "secondary" computer system.

At any given time, the computer system that is designated as the primary computer system performs all operations, that is, all transactions, of the system and generates a log of all such transactions. The computer system that is designated as the secondary computer system does not perform any operations while the primary system is in operation, but stores a copy of the transaction log of the primary system.

If there is a failure in the current primary computer system, the secondary computer system becomes the primary computer system and assumes execution of all transactions directed to the system, starting with the copy of the transaction log stored in the secondary computer system, which then becomes the primary computer system. The object of storing a

copy of the current primary system transaction log in the secondary system is to maintain the primary and secondary systems in "synchronization" so that the secondary computer system can assume execution of the transactions directed to the system with minimum loss of the preceding transactions that have been executed by the computer system that was previously the primary computer system.

For this purpose, the primary and secondary computer systems each stores a record of the "state" of the computer systems wherein the state of each computer system indicates whether the computer system is the current primary computer system or is the current secondary computer system. A third bit of "state" information is whether or not the two computer systems are in "synchronization", that is, have matching copies of the primary computer system transaction log. At any given time only one computer system can be the current primary computer system and can execute the system transactions, and the other computer system will always be the current secondary computer system and will be idle except for storing a copy of the primary computer system transaction log while waiting to assume execution of the transactions upon failure of the current primary computer system.

Lastly, the primary system transmits a copy of the primary system transaction records to the secondary system each time the transaction log in the primary system is "flushed" to disk, that is, is moved from the primary system memory space to the primary system mass storage device for long term storage. The primary system will flush its local copy of the transaction log to its own disk either when it transmits the copy of the transaction log to the secondary system or when it has transmitted the copy of the transaction log to the secondary system and has received an acknowledgment of the transmission from the secondary system.

It is, therefore, apparent that the present invention is distinguished over and from the Rastogi et al. '449 system for a number of fundamental reasons, which are recited in the claims as amended herein.

For example, the primary and secondary computer systems of the Rastogi et al. '449 system do not correspond either structurally or functionally with the dual control/processing

sub-systems of the present invention. That is, the two computer systems of the Rastogi et al. `449 system are not parallel, cooperating sub-systems within a system, but are completely and separate computer systems.

In still further fundamental distinction from the present invention, the dual control/processing sub-systems of the present invention operate independently but concurrently with each control/processing sub-system executing only the requests for transactions or operations that are directed to it, so that the full processing power of the two control/processing sub-systems is typically available at all times to process requests from clients. In basic contrast from the present invention, the two computer systems of the Rastogi et al. `449 system do not operate concurrently at any time. Instead, the primary processor operates at all times to execute all transactions of the system while the secondary processor is always idle with respect to the execution of transactions, so that only one half the potential processing power of the Rastogi et al. `449 system is available at any time.

In addition, the primary and secondary computer systems of the Rastogi et al. `449 system do not correspond either structurally or functionally with the state machine logging mechanism and state machine log mirroring mechanism of the present invention as each of the primary and secondary computer systems of the Rastogi et al. `449 system are a full function, general purpose computer capable of performing both transaction operations and transaction logging. In contrast, the state machine logging mechanism and state machine log mirroring mechanism of the present invention are both dedicated purpose, specialized function mechanisms that are structurally and functionally different from one another and are directed to separate and distinctly different functions. In a like manner, the control/processing sub-system and a corresponding state machine log mirroring mechanism are structurally and functionally distinguished from the primary and secondary computer systems of the Rastogi et al. `449 system for the same reasons.

It must also be noted that a control/processing sub-system and its associated state machine logging mechanism with the associated state machine log mirroring mechanism

cannot be compared, structurally or functionally, with the primary and secondary computer systems of the Rastogi et al. `449 system because the primary and secondary computer systems of the Rastogi et al. `449 system are in fact identical but completely separate and independent systems from one another. In contrast, the state machine log mirroring mechanism is functionally an integral element of the corresponding state machine logging mechanism, even through the state machine log mirroring system resides separately from the state machine logging mechanism so as not be involved in a failure of the corresponding control/processing sub-system with which the state machine log generator and log reside.

In further fundamental distinction between the present invention as recited in the claims and the Rastogi et al. `449 system, it must be noted that the state machine logging system and state machine log mirroring mechanism of the present invention generate or capture log state machine information, that is, state machines, and may do so at several points during any transaction or operation. The state machine logging system and state machine log mirroring mechanism of the present invention thereby generate or capture several state machine information records during the execution of a transaction or operation, and generate or capture each state machine concurrently with and during the execution of the transaction of operation.

In fundamental contrast from the present invention, the Rastogi et al. `449 system captures or generates only transaction records, that is, a record of a completed transaction at the conclusion of the transaction, and actually records, or stores, a transaction record only when the record is transferred into mass storage. In a like manner, the copy of a transaction record is transmitted to and stored in the secondary computer system only when the record is transferred to the primary system mass storage.

In still further fundamental distinction between the present invention and the Rastogi et al. `449 system, it must be noted that the system of the present invention captures and stores state information and the state machines represented by state information. As described herein above, and according to the present invention, a "state machine" is defined and is comprised of "machine state information", which is the control and data values residing

in a machine, mechanism or system at a given point in time and which represent the current operating state of the system and the operations then being performed by the system.

In fundamental contrast from the present invention, the Rastogi et al. `449 system captures or generates only transaction records, that is, the data comprising a record of a completed transaction, which has no relationship to control and data values representing the operating state of a system of operation at a given time. The difference is analogous to the difference between an entry in a checkbook and an operation by operation record of the control and data values in an accounting computer system.

In this regard, it must be noted that while the Rastogi et al. `449 system refers to computer system "state", this refers only to the operational status of the two systems and, more specifically, as to which system is the primary computer system and which system is the secondary computer system. The term "state" as used in the Rastogi et al. `449 system refers to system status, rather than to the operating states within the systems during the execution of operations by the systems.

It is the belief and position of the Applicant that for the reasons discussed above the teachings of Rastogi et al. `449 do not and cannot describe or suggest the present invention as recited in the claims under either or both of 35 U.S.C. § 102 or 35 U.S.C. § 103.

In this regard, it will be noted that the Applicant has amended the claims herein above to more explicitly and clearly recited these fundamental distinctions between the present invention and the prior art cited by the Examiner. These amendments, however, do not add any new matter to the present application or to the claims, and do not alter the subject matter or scope of the claims.

The Applicant, therefore, respectfully requests that the Examiner reconsider and withdraw all rejections of the claims as amended herein under either or both of 35 U.S.C. § 102 or 35 U.S.C. § 103 over Rastogi et al. `449, and the allowance of the claims as amended herein.

If any further amendment to this application is believed necessary to advance prosecution and place this case in allowable form, the Examiner is courteously solicited to contact the undersigned representative of the Applicant to discuss the same.

In view of the foregoing, it is respectfully submitted that the raised rejection(s) should be withdrawn and this application is now placed in a condition for allowance. Action to that end, in the form of an early Notice of Allowance, is courteously solicited by the Applicant at this time.

The Applicant respectfully requests that any outstanding objection(s) or requirement(s), as to the form of this application, be held in abeyance until allowable subject matter is indicated for this case.

In the event that there are any fee deficiencies or additional fees are payable, please charge the same or credit any overpayment to our Deposit Account (Account No. 04-0213).

Respectfully submitted,

Gary D. Clapp, Reg. No. 29,055
Customer No. 020210
Davis & Bujold, P.L.L.C.
Fourth Floor
500 North Commercial Street
Manchester NH 03101-1151
Telephone 603-624-9220
Facsimile 603-624-9229
E-mail: patent@davisandbujold.com

CERTIFICATE OF TRANSMISSION

I hereby certify that this correspondence is being transmitted via facsimile to the United States Patent and Trademark Office on: July 23, 2003

Gary D. Clapp

Type name of person signing certificate